

Social Welfare in One-sided Matching Markets without Money

Anand Bhalgat*

Deeparnab Chakrabarty*

Sanjeev Khanna*

Abstract

We study social welfare in one-sided matching markets where the goal is to efficiently allocate n items to n agents that each have a complete, private preference list and a unit demand over the items. Our focus is on allocation mechanisms that do not involve any monetary payments. We consider two natural measures of social welfare: the *ordinal welfare factor* which measures the number of agents that are at least as happy as in some unknown, arbitrary benchmark allocation, and the *linear welfare factor* which assumes an agent's utility linearly decreases down his preference lists, and measures the total utility to that achieved by an optimal allocation.

We analyze two matching mechanisms which have been extensively studied by economists. The first mechanism is the random serial dictatorship (RSD) where agents are ordered in accordance with a randomly chosen permutation, and are successively allocated their best choice among the unallocated items. The second mechanism is the probabilistic serial (PS) mechanism of Bogomolnaia and Moulin [8], which computes a fractional allocation that can be expressed as a convex combination of integral allocations. The welfare factor of a mechanism is the infimum over all instances. For RSD, we show that the ordinal welfare factor is asymptotically $1/2$, while the linear welfare factor lies in the interval $[\frac{526}{1000}, \frac{2}{3}]$. For PS, we show that the ordinal welfare factor is also $1/2$ while the linear welfare factor is roughly $2/3$. To our knowledge, these results are the first non-trivial performance guarantees for these natural mechanisms.

*Dept. of Comp. and Inf. Science, Univ. of Pennsylvania, Philadelphia, PA 19104.
Email: bhalgat,deepc@seas.upenn.edu, sanjeev@cis.upenn.edu

1 Introduction

In the one-sided matching market problem¹, the goal is to efficiently allocate n items, I , to n unit-demand agents A , with each agent a having a complete and private preference list \geq_a over these items. The problem arises in various applications such as assigning dormitory rooms to students, time slots to users of a common machine, organ allocation markets, and so on. Since the preferences are private, we focus on *truthful* (strategyproof) mechanisms in which agents do not have an incentive to misrepresent their preferences. One class of such mechanisms involve monetary compensations/payments among agents. However, in many cases (e.g., in the examples cited above), monetary transfer may be infeasible due to reasons varying from legal restrictions to plain inconvenience. Hence, we focus on truthful mechanisms without money.

A simple mechanism for the one-sided matching problem is the following: agents arrive one-by-one according to a fixed order σ picking up their most preferred unallocated item. This is called a *serial dictatorship* mechanism. The *random serial dictatorship* (RSD) mechanism picks the order σ uniformly at random among all permutations. Apart from being simple and easy to implement, RSD has attractive properties: it is truthful, fair, anonymous/neutral, non-bossy², and returns a Pareto optimal allocation. In fact, it is the only truthful mechanism with the above properties [31], and there is a large body of economic literature on this mechanism (see Section 1.2).

Despite this, an important question has been left unaddressed: how *efficient* is this mechanism? To be precise, what is the guarantee one can give on the social welfare obtained by this algorithm when compared to the optimal social welfare? As computer scientists, we find this a natural and important question, and we address it in this paper.

How should one measure the social welfare of a mechanism? The usual recourse is to *assume* the existence of *cardinal* utilities u_{ij} of agent i for item j which induces the preference list over items with the semantic that agent i prefers item j to ℓ iff $u_{ij} > u_{i\ell}$. A mechanism has *welfare factor* α if for every instance the utility of the matching returned is at least α times that of the optimum utility matching. There are a couple of issues with this. Firstly, it is not hard to see that one cannot say anything meaningful about the performance of RSD if the utilities are allowed to be arbitrary. This is because the optimum utility matching might be arising due to one particular agent getting one particular item (a single edge), however with high probability, any random permutation would lead to another agent getting the item and lowering the total welfare by a lot³. Secondly, the assumption of cardinal utilities inherently ties up the performance of the algorithm with the ‘cardinal numbers’ involved; the very quantities whose existence is only an assumption. Rather, what is needed is an *ordinal* scale of analyzing the quality of a mechanism; a measure that depends only on the order/preference lists of the agents rather than the precise utility values.

In this paper, we propose such a measure which we call the *ordinal social welfare* of a mechanism. Given an instance of items and agents with their preference lists, we assume that there exists some benchmark matching M^* , unknown to the mechanism. We stress here this can be *any* matching. We say that the *ordinal welfare factor* of a mechanism is α , if for any instance and every benchmark matching M^* , in the matching returned by the mechanism, at least αn agents (in expectation) get an item which they prefer at least as much as their allocation in M^* .

¹In the literature, the problem has been alternately called the *house allocation* or *assignment* problem.

²A mechanism is neutral if the allocation of items doesn’t change with renaming, and is non-bossy if no agent can change his preference so that his allocation remains unchanged while someone else’s changes.

³The reader may notice similarities of RSD with online algorithms for bipartite matching problems. We elaborate on the connection in Section 2.2.

A discussion of this measure is in order. Firstly, the measure is *ordinal* and is well defined whenever the utilities are expressed via preference lists. Secondly, the notion is independent of any ‘objective function’ that an application might give rise to since it measures the ordinal social welfare with respect to any desired matching. One disadvantage of the concept is that it is global: it counts the fraction of the *total* population which gets better than their optimal match. In other words, if everyone is “happy” in the benchmark matching M^* , then a mechanism with the ordinal welfare factor α will make an α fraction of the agents happy. However if M^* itself is *inefficient*, say only 1% of the agents are “satisfied” in M^* , then the ordinal welfare factor does not say much. For instance, it does not help for measures like “maximize number of agents getting their first choice”, for in some instances, this number could be tiny even in M^* . Furthermore, it does not say anything about the “happiness” of individual agents, e.g. a mechanism may have the ordinal welfare factor close to 1, but there may exist an agent who is almost always allocated an item that he prefers less than M^* . Finally, we observe that the ordinal welfare factor of any mechanism, even ones which know the true preference lists, cannot be larger than $1/2$. The reason for this is that the allocation must be competitive with respect to all benchmark matchings simultaneously, and it can be seen (Theorem 7) that in the instance when all agents have the same preference list, if M^* is chosen to be a random allocation, then no mechanism can have an ordinal welfare factor better than $1/2$. As our first result, we show that the ordinal welfare factor of RSD is in fact asymptotically $1/2$.

Theorem 1. *Given any instance and any matching M^* , the expected fraction of agents getting an item in RSD which is as good as what they get in M^* is at least $1/2 - o(1)$.*

Till now we have focussed on the RSD mechanism since it is a simple (and almost unique) truthful mechanism for the matching market problem. A mechanism is called truthful if misrepresenting his preference list doesn’t strictly increase the total *utility* of an agent, where the utility is defined as the cardinal utility obtained by the agent on getting his allocated item. However, when the utilities of agents are represented as preference lists, one needs a different definition. In light of this, Bogomolnaia and Moulin [8] proposed a notion of truthfulness based on the *stochastic dominance*. Roughly speaking, for an agent a random allocation rule stochastically dominates another if the probability of getting one of his top k choices in the first rule is at least that in the second, for any k .

Using this notion, [8] called a mechanism (weakly) truthful if no agent can obtain a stochastically dominating allocation by misreporting his preference list. In other words, there exists at least one utility function u that maps a preference list over items to cardinal utility values, such that the mechanism is truthful if the user’s utility function is u . With this definition, the authors propose a remarkable mechanism, which they called the *probabilistic serial* (PS) algorithm, and prove that it is truthful in the above sense. At a high level, instead of allocating “best” items integrally to agents (as RSD does), the PS algorithm does so fractionally obtaining a fractional perfect matching. By the Birkhoff-von Neumann decomposition theorem in combinatorial optimization, this implies a distribution on integral matchings, and this is the output of the PS algorithm. We give a complete description of the PS algorithm in Section 1.1.

PS and RSD are incomparable and results on RSD do not a priori imply those for PS. There exist instances where RSD does strictly better than PS, and there exist instances where vice-versa is true (see Appendix A). Nevertheless, PS has an ordinal welfare factor of at least $1/2$ as well.

Theorem 2. *Given any instance and matching M^* , the expected fraction of agents which get an item in PS which is as good as their allocation in M^* , is at least $1/2$.*

Ordinal Welfare Factor and Popular Matchings Our notion of ordinal welfare factor is somewhat related to the notion of *popular* matchings [16, 4, 24]. Given preference lists of agents, a matching M is said to be more popular than M' if the number of agents getting *strictly* better items in M is at least the number of agents getting strictly better items in M' . A matching is popular if no other matching is more popular than it. Thus the notion of popularity partitions the set of agents into three groups while comparing two matching M and M' – (a) agents that prefer M , (b) agents that are neutral, and (c) agents that prefer M' over M , where as in the case of ordinal welfare factor, agents that prefer M , and agents that are neutral, are not distinguished.

It can be easily seen that any popular matching has an ordinal welfare factor of at least $1/2$, however, (a) not every input instance has a popular matching, and (b) no truthful algorithms are known to compute them when they exist. A few modified measures for matchings have been studied in the literature [25, 20, 24], namely (a) unpopularity factor ($\#$ agents that are unhappy/ $\#$ agents that are strictly happier), and (b) unpopularity margin ($\#$ agents that are unhappy – $\#$ agents that are strictly happier), (c) popular mixed matchings. The ordinal welfare factor can be seen as a new way of quantifying the popularity of a matching.

Linear Utilities. We also analyze the performance of RSD and PS mechanisms when agents' utilities are linear, one of the most commonly studied special case of cardinal utilities. In this model, we assume that the utility for an agent for his i^{th} preference is $\frac{n-i+1}{n}$.

A quick observation shows that *any* serial dictatorship mechanism achieves a welfare of at least 50% of the optimum: the agent at step t gets his t^{th} choice or better giving him a utility of at least $(1 - (t - 1)/n)$; thus the total utility is at least $(n + 1)/2$, and the optimum is at most n . How much better does RSD do? Intuitively, from the above calculation, one would expect the worst case instance would be one where each agent gets one of his top $o(n)$ choices; that would make the optimum value $(1 - o(1))n$. We call such instances as *efficient* instances since there is an optimum matching where every one gets their (almost) best choice. We show that for efficient instances, RSD achieves (asymptotically) at least $2/3$ of the social optimum. Furthermore, there exists instances where RSD does no better. These bounds hold for PS mechanism as well.

Theorem 3. *When the utilities are linear and the instance is efficient, the welfare factor of RSD is at least $2/3 - o(1)$. Furthermore, there exist efficient instances for which the welfare factor of RSD is at most $2/3 + o(1)$.*

Theorem 4. *When the utilities are linear and the instance is efficient, the welfare factor of PS is at least $2/3 - o(1)$. Furthermore, there exist efficient instances for which the welfare factor of PS is at most $2/3 + o(1)$.*

For general instances, we can show that PS achieves a social welfare of at least 66% (but not quite $2/3$) of the optimum (see Appendix C for details). However, we cannot extend our intuition and prove a similar performance for RSD on general instances; a ‘balancing’ argument shows that the performance of RSD for linear utilities is at least 52.6%, although we conjecture this can be improved (see Appendix B for details).

Extensions. We consider two extensions to our model and focus on the performance of RSD leaving that of PS as an open direction. In the first, we let the preference lists be incomplete. The proof of Theorem 1 implies that the ordinal welfare factor of RSD remains unchanged. For linear utilities, we generalize the definition as follows: for an agent with a preference list of length ℓ , the first choice gives him a utility of 1 while the last item gives a utility of $1/\ell$. Our results show that RSD doesn't perform very well in this case.

Theorem 5. *For linear utilities, RSD gets at least $\tilde{\Omega}(n^{-1/3})$ fraction of the social optimum. Furthermore, there are instances, where the welfare of RSD is at most $\tilde{O}(n^{-1/3})$ fraction of the social optimum.*

In the second extension, we let the demand of an agent be for sets of size K or less, for some $K \geq 1$. Agents now arrive and pick their best ‘bundle’ among the unallocated items. The ordinal welfare factor of a mechanism is now α if at least an α fraction of agents get a bundle that is as good (assuming there is a complete order on the set of bundles) as what they got in an arbitrary benchmark allocation. We show that RSD has ordinal welfare factor $\Theta(1/K)$.

Theorem 6. *In the case when each agent has a maximum demand of K items, the ordinal welfare factor of RSD is $\Theta(1/K)$.*

1.1 Preliminaries

Utility Models, Truthful Mechanisms, Welfare Factors. As stated above, we consider two models for utilities of agents. In the *cardinal utility model*, each agent a has a utility function $u_a : I \rightarrow \mathbb{R}_{\geq 0}$, with the property that $j >_a \ell$ iff $u_a(j) > u_a(\ell)$. Given a distribution on the matchings, the utility of agent a is $u_a(M) := \sum_{M \in \mathcal{M}} p(M) u_a(M(a))$, where $p(M)$ is the probability of matching M . In this paper, we focus on the special case of *linear utility model* where the i^{th} ranked item for any agent a gives him a utility of $(1 - (i - 1)/n)$. We call an instance *efficient*, if there is a matching which matched every agent to an item in his top $o(n)$ (for concreteness, let’s say this is $n^{1/5}$) choices and thus gives him utility of $(1 - o(1))$. In the *ordinal utility model*, each agent a represents his utility only via his complete preference list \geq_a over the items.

A mechanism \mathcal{A} is *truthful* if no agent can misrepresent his preference and obtain a better item. In the cardinal utility model this implies that for all agents a and utility functions u_a, u'_a , we have $u_a(M) \geq u_a(M')$ where $M = \mathcal{A}(u_1, \dots, u_n)$ and $M' = \mathcal{A}(u_1, \dots, u'_a, \dots, u_n)$. In the ordinal utility model, following [8], we call a mechanism \mathcal{A} to be truthful if for all agents a the following holds. Let $M = \mathcal{A}(>_1, \dots, >_n)$, and let $M' = \mathcal{A}(>_1, \dots, >'_a, \dots, >_n)$ be the solution when a changes his preference list. Then the randomized allocation to a returned by M' *does not* stochastically dominate that returned by M . In fact, [8] define a stronger notion of truthfulness in which the allocation to a returned by M stochastically dominates that returned by M' . This is a strong notion of truthfulness since it implies truthfulness for *any* cardinal utilities as well – if M stochastically dominates M' , then for any utility function $u_a(M) \geq u_a(M')$.

A mechanism has *linear welfare factor* of α if for all instances the (expected) sum of linear utilities of agents obtained from the allocation of the mechanism is at least α times the optimal utility allocation for that instance. A mechanism has *ordinal welfare factor* of α if for all instances, and for *all benchmark matchings* M^* , at least α fraction of agents (in expectation) obtain an item via the mechanism which is at least as good as that allocated in M^* .

The Probabilistic Serial Mechanism. The *probabilistic serial* (PS) mechanism was suggested by Bogomolnaia and Moulin [8]. The mechanism fractionally allocates items to agents over multiple phases, we denote the fraction of the item i allocated to an agent a by $x(a, i)$. These fractions are such that $\sum_{a \in A} x(a, i) = \sum_{i \in I} x(a, i) = 1$ for all agents a and items i . Thus this fractional allocation defines a distribution on integral matchings, and this is the distribution returned by PS.

Initially, $x(a, i) = 0$ for every agent a and item i . We say that an item i is allocated if $\sum_{a \in A} x(a, i) = 1$, otherwise we call it to be available. The algorithm grows $x(a, i)$ ’s in phases, and in each phase one or more items get completely allocated. During a phase of the algorithm, each

agent a grows $x(a, i)$ at the rate of 1 where i is his best choice in the set of available items. The current phase completes and the new phase starts when at least one item that was available in the current phase, gets completely allocated. The algorithm continues until all items are allocated.

We make a few observations about the above algorithm which will be useful in our analysis. Firstly note that the algorithm terminates at time $t = 1$, at which time all agents are fractionally allocated one item, that is, $\sum_{i \in I} x(a, i) = 1$. This is because the LHS grows at the rate of 1 for all agents in any point of time (one uses the fact the preference lists are complete). Secondly, note that any phase lasts for time at least $1/n$ and at most 1. Therefore, by time $< j/n$ for any $1 \leq j \leq n$, at most $(j - 1)$ phases would've completed.

1.2 Related Work

As stated in the introduction, there is a huge amount of literature on matching markets starting with the seminal paper of Gale and Shapley [15], and we refer the reader to numerous detailed surveys (the classic Roth and Sotomayor [28], and more recent ones by Sönmez and Ünver [30] and Abdulkadiroglu and Sönmez [3]) for a more elaborate picture. In this section, we try to review the papers which are most relevant to our paper.

The one-sided matching market design problem was first studied by Hylland and Zeckhauser [21] who propose a mechanism to find a distribution on matchings via a market mechanism where each agent is given equal, artificial budgets. Their mechanism returns Pareto optimal, envy-free solutions; unfortunately it is not truthful. Zhou [32], answering a question of Gale [14], showed that there can be no truthful mechanism which is anonymous/neutral and satisfies *ex ante* Pareto optimality; this is a stronger notion of *ex post* Pareto optimality which we have discussed before. Svensson [31] showed that serial dictatorship mechanisms are the only truthful mechanisms which are (ex post) Pareto optimal, non bossy, and anonymous. The study of mechanisms with ordinal utilities for this problem was started by Bogomolnaia and Moulin[8]. The probabilistic serial mechanism was in fact proposed in an earlier paper on scheduling jobs by Cres and Moulin [11]. Following the work of [8], there was a list of work characterizing stochastic dominance [2, 9], and generalizing it to the case of incomplete preference lists [23], and to multiple copies of items [10].

A generalization of the one-sided matching market design problem where the items are initially owned by agents, and they need to be exchanged among themselves is called the *market exchange* problem (or the *housing market* problem.) This problem has many applications, the most important of which is the kidney exchange problem [27]. The exchange problem was introduced by Scarf and Shapley [29] who were interested in *core* allocations; allocations where no subset of agents can re-allocate their items and everyone be better off. They showed the core allocation is non-empty via a constructive algorithm attributed to David Gale; this is Gale's *top trading cycle* (TTC) algorithm which is truthful and Pareto optimal. There has been a lot of work extending TTC to various scenarios and characterizing truthful mechanisms in this model (we refer the interested reader to the surveys [30] and [3] and the references therein); however the result most relevant to work is the one due to Abdulkadiroglu and Sönmez [1]. They showed that in the one-sided market mechanism problem, if one assumes a random initial endowment and runs the TTC algorithm, then the behaviour (distribution on matchings) is exactly the same as that of RSD.

Finally, the study of mechanism design without money has been of recent interest in the computer science community [26, 12]. We already have mentioned the relation to popular matchings in the introduction. There has been works motivated by the market exchange problem [6, 13], and item allocation problem [18, 19], however none of them address the problem that we study.

2 Ordinal Welfare Factor of RSD and PS Mechanisms

In this section, we prove Theorems 1 and 2. We first show that the ordinal welfare factor of *any* mechanism is at most $1/2$ in the instance where every agent has the same preference list.

Theorem 7. *If every agent has the same preference list $(1, 2, \dots, n)$, then the ordinal welfare factor of any mechanism is at most $1/2 + 1/2n$.*

Proof. A mechanism returns a probability distribution on matchings which we will interpret as a distribution of permutations, let \mathcal{D} be that distribution. We choose the benchmark matching M^* to be a random matching, i.e. a random permutation π of the agents.

It suffices to show that for any fixed permutation $\sigma \in \mathcal{D}$, the expected number of agents a such that $\sigma(a) > \pi(a)$ is $(n-1)/2$. The bound then holds for \mathcal{D} . Since π is chosen uniformly at random, the probability that $\pi(a) \leq \sigma(a)$ is precisely $\sigma(a)/n$. So the expected number of happy people for the permutation σ is $\sum_{a \in A} \sigma(a)/n = (n+1)/2$. \square

2.1 Ordinal Welfare Factor of RSD

In this section, we prove Theorem 1. Let M^* be the unknown benchmark matching. We call an agent a *dead* at time t if he hasn't arrived yet and all items as good as $M^*(a)$ in his preference list has been allocated. Let D_t be the expected number of dead agents at time t . Let ALG_t be the expected number of agents who get an item as good as their choice in M^* by time t . From the above definition, we get

$$\text{ALG}_{t+1} - \text{ALG}_t = 1 - \frac{D_t}{n-t} \quad (1)$$

We will now bound D_t from above which along with (1) will prove the theorem.

Lemma 1. $D_t \leq \frac{(t+2)(n-t)}{n+1}$ for $1 \leq t \leq n$.

Before proving the lemma, note that adding (1) for $t = 1$ to $n-1$ gives $\text{ALG}_n - \text{ALG}_1 \geq \sum_{t=1}^{n-1} \left(1 - \frac{t+2}{n+1}\right)$, implying $\text{ALG}_n - \text{ALG}_1 \geq n/2 - 2n/n$. This proves that the ordinal welfare factor of RSD is at least $1/2 - o(1)$ proving Theorem 1.

Proof. Let us start with a few definitions. For an item i and time t , let $\text{ALL}_{i,t}$ be the event that item i is allocated by time t . For an agent a and time t , let $\text{LATE}_{a,t}$ be the event that a arrives after time t . The first observation is this: if an agent a is dead at time t , then the event $\text{ALL}_{M(a),t}$ and $\text{LATE}_{a,t}$ must have occurred. Therefore we get

$$D_t \leq \sum_{a \in A} \Pr[\text{ALL}_{M(a),t} \wedge \text{LATE}_{a,t}] \quad (2)$$

Note that $\Pr[\text{LATE}_{a,t}]$ is precisely $(1 - t/n)$. Also, note that $\sum_{i \in I} \Pr[\text{ALL}_{i,t}] = t$. This is because all agents are allocated *some* item. Now suppose *incorrectly* that $\text{ALL}_{M(a),t}$ and $\text{LATE}_{a,t}$ were independent. Then, (2) would give us

$$D_t \leq \left(1 - \frac{t}{n}\right) \sum_{a \in A} \Pr[\text{ALL}_{M(a),t}] = \left(1 - \frac{t}{n}\right) \sum_{i \in I} \Pr[\text{ALL}_{i,t}] = \frac{t(n-t)}{n} \quad (3)$$

which is at most the RHS in the lemma. However, the events are not independent, and one can construct examples where the above bound is indeed incorrect. To get the correct bound, we need the following claim.

Claim 1.

$$\frac{\Pr[\text{ALL}_{M(a),t} \wedge \text{LATE}_{a,t}]}{(n-t)} \leq \frac{\Pr[\text{ALL}_{M(a),t+1} \wedge \overline{\text{LATE}_{a,t+1}}]}{(t+1)}$$

Proof. This follows from a simple charging argument. Fix a relative order of all agents other than a and consider the n orders obtained by placing a in the n possible positions. Observe that if the event $\text{ALL}_{M(a),t} \wedge \text{LATE}_{a,t}$ occurs at all, it occurs exactly $(n-t)$ times when a 's position is $t+1$ to n . Furthermore, crucially observe that if the position of a is 1 to $t+1$, the item $M(a)$ will still be allocated. This is because the addition of a only leads to worse choices for agents following him and so if $M(a)$ was allocated before, it is allocated even now. This proves that for every $(n-t)$ occurrences of $\text{ALL}_{M(a),t} \wedge \text{LATE}_{a,t}$, we have $(t+1)$ occurrences of the event $\text{ALL}_{M(a),t+1} \wedge \overline{\text{LATE}_{a,t+1}}$. The claim follows as it holds for every fixed relative order of other agents. \square

Now we can finish the proof of the lemma. From Claim 1, we get

$$\frac{t+1}{n-t} \cdot \Pr[\text{ALL}_{M(a),t} \wedge \text{LATE}_{a,t}] \leq \Pr[\text{ALL}_{M(a),t+1}] - \Pr[\text{ALL}_{M(a),t+1} \wedge \text{LATE}_{a,t+1}]$$

Taking the second term of the RHS to the LHS, adding over all agents, and invoking (2), we get

$$\frac{t+1}{n-t} \cdot D_t + D_{t+1} \leq t+1 \tag{4}$$

Using the fact that $D_{t+1} \geq D_t - 1$ (the number of dead guys cannot decrease by more than 1), and rearranging, proves the lemma. \square

2.2 RSD and online bipartite matching

In this section, we highlight the relation between RSD and algorithms for online bipartite matching. As we show below, the analysis of RSD above can be seen as a generalization of online bipartite matching algorithms. This section can be skipped without loss of continuity.

In the online bipartite matching problem, vertices of one partition (think of them as agents) are fixed while vertices of the other partition (think of them as items) arrive in an adversarial order. When an item arrives, we get to see the incident edges on agents. These edges indicate the set of agents that desire this item. The algorithm must immediately match this item to one of the unmatched agents desiring it (or choose to throw it away). In the end, the size of the obtained matching is compared with the optimum matching in the realized graph.

Karp, Vazirani and Vazirani [22] gave the following algorithm (KVV) for the problem: fix a *random* ordering of the agents, and when an item arrives give it to the first unmatched agent in this order. They proved⁴ that the expected size of the matching obtained is at least $(1 - 1/e)$ times the optimum matching. For simplicity (and indeed this can be shown to be without loss of generality), assume that the size of the optimum matching is indeed n , the number of agents and items. The KVV theorem can be ‘flipped around’ to say the following. Suppose each agent has the preference list which goes down its desired items in the order of entry of items. Then, if agents arrive in a random order and pick their best, unallocated, desired item, in expectation an $(1 - 1/e)$ fraction of agents are matched. That is, if we run RSD on this instance (with incomplete lists), an $(1 - 1/e)$ fraction of agents will get an item.

⁴In 2008, a bug was found in the original extended abstract of [22], but was soon resolved. See [17, 7, 5] for discussions and resolutions.

We should point out that the above result does not a priori imply an analysis of RSD, the reason being that in our problem an agent a , when he arrives, is allocated an item even if that item is *worse* than what he gets in the benchmark matching M^* . This might be bad since the allocated item could be ‘good’ item for agents to come. In particular, if the order chosen is not random but arbitrary, the performance of the algorithm is quite bad; in contrast, the online matching algorithm still has a competitive ratio of $1/2$. Concretely, consider the following instance: Agent i for $2 \leq i \leq n$ have item $(i-1)$ as their first choice, and item i as their second choice. Agent 1 has item 1 as his first choice and item n as his second choice. M^* matches agent $2, \dots, n$ to items $1, \dots, (n-1)$ and agent 1 to item n . Now if agents arrive in order 1 to n , only agent 1 does as well as its allocation in M^* . We now tailor our analysis of RSD to give another proof of the online bipartite matching result.

Corollary 1. *KVV matches at least $(1 - 1/e)n - 1$ agents in expectation.*

Proof. The analysis is exactly same as the proof of Theorem 1 except we have the following ‘savings’ – instead of $\sum_{i \in I} \Pr[\text{ALL}_{i,t}] = t$, we have $\sum_{i \in I} \Pr[\text{ALL}_{i,t}] = \text{ALG}_t$. This is because by time t the number of items allocated is precisely the number of agents who receive an item, which is ALG_t . (4) then evaluates to $\frac{t+1}{n-t} \cdot D_t + D_{t+1} \leq \text{ALG}_{t+1}$. Using the fact that $D_{t+1} \geq D_t - 1$ and $\text{ALG}_{t+1} \leq \text{ALG}_t + 1$, we get $\frac{1}{n-t} \cdot D_t \leq \frac{1}{n+1} \cdot \text{ALG}_t + \frac{2}{n+1}$, which along with (1) gives us

$$\text{ALG}_{t+1} \geq \left(1 - \frac{1}{n+1}\right) \text{ALG}_t + \left(1 - \frac{2}{n+1}\right)$$

This solves to $\text{ALG}_n \geq (1 - 1/e) \cdot (n - 1)$. □

2.3 Ordinal Welfare Factor of PS

In this section, we prove Theorem 2. We suggest the reader to refer to the algorithm and its properties as described in Section 1.1. In particular, we will use the following observation.

Observation 1: By time $< j/n$, for any $1 \leq j \leq n$, at most $(j-1)$ items are completely allocated.

Let M^* be the unknown benchmark matching. For an agent a , let t_a be the time in the run of the PS algorithm at which the item $M^*(a)$ is completely allocated. Observe that the probability agent a gets an item $M^*(a)$ or better is precisely t_a , since till this time $x(a, i)$ increases for items $i \geq_a M^*(a)$. Summing up all agents, we see that the ordinal welfare factor of the PS mechanism is $\sum_a t_a$. The observation above implies at most $(j-1)$ agents have $t_a < j/n$. So, $\sum_a t_a \geq \sum_{j=1}^n (n - j + 1)/n \geq n/2 + 1/2$. This completes the proof of Theorem 2.

3 Linear Welfare Factor of RSD and PS

In this section, we establish bounds on the linear welfare factor of RSD and PS mechanisms. We first prove Theorem 3 in two lemmas. Recall that an instance is called efficient if there exists a matching in which every agent is matched to an item in his top $o(n)$ choices.

Lemma 2. *When the utilities are linear and the instance is efficient, the linear welfare factor of RSD is at least $(2/3 - o(1))$.*

Proof. The proof almost directly follows from Lemma 1. Let U_t denote the expected utility obtained by time t . Consider the agent coming at time $t+1$. If he is not dead already, then he will get a

utility of at least $(1 - o(1))$ (since the instance is efficient). If he is dead, then he will get a utility of at least $(1 - t/n)$. This is because only t items have been allocated and this agent takes an item $(t + 1)$ th ranked or higher. Therefore,

$$U_{t+1} - U_t \geq \left(1 - \frac{D_t}{n-t}\right) \cdot (1 - o(1)) + \frac{D_t}{n-t} \cdot (1 - t/n) \geq 1 - o(1) - \frac{t}{n} \cdot \frac{D_t}{n-t}$$

Using Lemma 1, we get $U_{t+1} - U_t \geq 1 - o(1) - \frac{t(t+2)}{n(n+1)}$. Summing over all t , we get that the total utility of RSD is at least $(1 - o(1))n - (n/3 + o(n)) = (2/3 - o(1))n$. \square

The above analysis can be modified via a balancing trick to give a strictly better than 50% guarantee for all instances in the case of linear utilities. However, the improvement we get is small, and the analysis is not that informative. We defer it to the Appendix B.

Lemma 3. *When the utilities are linear, there exists an efficient instance for which RSD gets a utility of at most $(2/3 + o(1))n$.*

Proof. Partition n agents and items into t blocks of size n/t each, where $t = n^{1/5}$. We denote the j^{th} block of agents and items by A_j and I_j respectively, and they number from $\left(\frac{(j-1)n}{t} + 1\right)$ to $\frac{jn}{t}$.

We now illustrate the preference lists of agents. Fix an agent a in block A_j . Let he be the k^{th} agent in the block, where $1 \leq k \leq n/t$, i.e. his agent number is $(j-1)n/t + k$. A random set of t^3 items is picked from each of blocks I_1, \dots, I_{j-1} , and these form the first $(j-1)t^3$ items in his preference list, in increasing order of item number. The item $(j-1)n/t + k$ is his $((j-1)t^3 + 1)^{th}$ choice. His remaining choices are the remaining items considered in increasing order. This completes the description of the preference lists of the agents.

Note that if every agent a is assigned the corresponding item with the same number, then each agent gets one of his top t^4 choices, leading to a utility of at least $(1 - \frac{t^4}{n}) = 1 - o(1)$. So, the instance is indeed efficient. We now show that RSD gets utility at most $2n/3 + o(1)$.

Let σ be a random permutation of the agents. We divide σ into t chunks of n/t agents, with the j^{th} chunk, S_j , consisting of agents $\sigma(\frac{(j-1)n}{t} + 1)$ to $\sigma(\frac{jn}{t})$. Note that with high probability ($\geq (1 - 1/t^3)$), we have that for any block A_j and chunk S_i , $|A_j \cap S_i| \in [(1 - \frac{1}{t^2})\frac{n}{t^2}, (1 + \frac{1}{t^2})\frac{n}{t^2}]$. We now state the crucial observation.

Claim 2. *With high probability, at least $(1 - \frac{1}{t^3})$ fraction of the items in the first i blocks have been allocated after arrival of first i chunks.*

Proof. Fix an agent a in the first i chunks and consider the allocation to the agent a when he arrives. If more than in/t^3 items in set of items in $I_1 \cup \dots \cup I_i$ are left unallocated, then we claim that the probability that the agent a is not allocated one of these free items is exponentially small. Note that it^3 items of the in/t items in the first i blocks form the top choices of agent a . Furthermore, this set is chosen *independently* of all other agents. In particular, the probability that none of the free items lie as one of these choices is at most $\left(1 - \frac{t^4}{n}\right)^{in/t^3} \leq e^{-t}$. Since t is large enough (assumed to be $n^{1/5}$), this probability is small enough to apply union bound over all agents. \square

Now we are ready to analyze RSD. Consider the $(i+1)^{th}$ chunk of agents. With high probability, there are at least $\frac{n}{t^2}(1 - \frac{1}{t^2})$ agents from each block A_1, \dots, A_i in S_{i+1} . Since only in/t^3 items remain from the first i block of items, at least $\frac{in}{t^2}(1 - \frac{1}{t^2}) - \frac{in}{t^3}$ of these agents must get an item from blocks

$(i + 1)$ or higher. However, this gives them utility at most $(1 - \frac{in/t}{n}) \geq 1 - i/t$. That is, the *drop* in their utility to what they get in the optimum is at least i/t . Summing the total drop over all agents and all chunks, we get that the difference between RSD and the optimum is at least

$$\sum_{i=1}^t \frac{in}{t^2} (1 - \frac{1}{t}) \frac{i}{t} = (1 - o(1)) \frac{n}{t^3} \sum_{i=1}^n i^2 = n/3$$

Therefore, the social welfare of RSD is at most $(2/3 + o(1))n$. \square

Linear Welfare Factor of PS Mechanism We now establish the upper and the lower bound on the linear welfare factor of the PS mechanism. As in the case of RSD, we focus on efficient instances in the main body since they contain the main analysis ideas.

Lemma 4. *When the utilities are linear and the instance is efficient, the linear welfare factor of the PS mechanism is at least $(2/3 - o(1))$.*

Proof. Let o_a denote the utility obtained by agent a in the utility optimal matching. Since the instance is efficient, $o_a = 1 - o(1)$ for all agents a .

Consider the j^{th} phase of PS, and suppose it lasts for time Δ_j . Observation 1 implies that $\sum_{j \leq \ell} \Delta_j \geq \ell/n$. Furthermore, in phase j , at least $(n - j + 1)$ agents are getting utility at a rate higher than their utility in the optimal matching. This is because at most $(j - 1)$ items have been completely allocated. Furthermore, the remaining $(j - 1)$ agents are getting utility at a rate at least $(1 - (j - 1)/n)$ since they are growing their $x(a, i)$ on their j^{th} choice or better. So, the total utility obtained by PS is at least $\sum_{j=1}^n \Delta_j \cdot \left((n - j + 1) \cdot (1 - o(1)) + (j - 1) \cdot (1 - \frac{j-1}{n}) \right) \approx \sum_{j=1}^n \Delta_j \left(\frac{n^2 - (j-1)^2}{n} \right)$

Now, the RHS is smallest if Δ_1 is as small as possible, modulo which, Δ_2 is as small as possible and so on. Given the constraint on Δ_j 's, we get that the RHS is at least $\sum_{j=1}^n \frac{n^2 - (j-1)^2}{n^2} \approx 2n/3$. This implies a $2/3$ approximation. \square

Lemma 5. *When the utilities are linear, there exists an efficient instance such that the total utility of PS is at most $(2/3 + o(1))$ times the optimum.*

Proof. The construction of the lower bound example is similar to (but not the same as) that given in the proof of Lemma 3. We have n agents and items divided in t blocks (we assume t divides n) of n/t agents each, t will be $o(n)$ in the construction. We denote the j^{th} block of agents and items by A_j and I_j respectively, and they number from $\left(\frac{(j-1)n}{t} + 1 \right)$ to $\frac{jn}{t}$.

The agents' preference lists are as follows. Fix an agent a in block A_j . Let he be the k^{th} agent in this block, where $1 \leq k \leq n/t$, i.e. his agent number is $(j - 1)n/t + k$. The first j preferences of this agent are precisely the k^{th} items in the blocks I_1 to I_j , in the increasing order of the item number. That is, for $1 \leq \ell \leq j - 1$, the $(\ell + 1)$ th choice of agent a is the item number $\ell n/t + k$. Furthermore, for $j < \ell \leq t$, the item number $\ell n/t + k$, present in block I_ℓ , is the $(\ell n/t + 1)^{th}$ choice of agent a . Thus, so far we have described where t items reside in the preference list of a . The remaining $n - t$ choices of his are all the remaining items from 1 to n , in that order.

Lemma 6. *The optimum social welfare of the instance is at least $n - t$.*

Proof. Note that agent a whose number is $(j-1)n/t + k$ can be assigned item $(j-1)n/t + k$, which is his j^{th} choice. So, in this matching, every agent gets his t^{th} choice or higher, leading to a social welfare of at least $n(1 - t/n)$. \square

Lemma 7. *In this instance, PS attains a social welfare of at most $(2/3 + o(1))n$.*

Proof. The following claim is the key observation in the proof of this lemma.

Claim 3. *In the run of PS, each phase is of duration $1/t$, and all items in set I_j are completely allocated in phase j .*

Note that initially all agents prefer items in I_1 , furthermore, each item in I_1 is the best choice of exactly t agents. So at time $1/t$, all items in I_1 are completely allocated. Subsequently, the best choice of every agent among the remaining items lies in I_2 and arguing inductively proves the claim. \square

Now we can complete the proof of the lemma. Note that each agent in block A_j obtains precisely $1/t$ fraction of an item from every block. For the first j blocks, he gets utility at most 1 per item. Subsequently, for $j < \ell \leq t$, he gets utility $(1 - \frac{(\ell-1)(n/t)}{n})$. So, the utility of an agent in A_j is at most $\frac{1}{t} \left(1 \cdot j + \sum_{\ell=j+1}^t \frac{t-\ell+1}{t} \right) \approx \frac{1}{t} \left(j + (t-j) - \frac{t^2-j^2}{2t} \right) = \frac{1}{2} + \frac{j^2}{2t^2}$ where the \approx hides the additive $o(1)$ term. Summing over all blocks j and all agents in a block we get that the social welfare of PS is at most $\frac{n}{2} + \frac{n}{6} + o(n) = 2n/3 + o(n)$. \square

For PS, we can do a much better analysis for general instances with linear utilities, than what we can do for RSD. We defer this to the Appendix C.

4 Extensions

In this section, we study two extensions to the model studied thus far, (a) incomplete preference list: where for every agent, there exists a set of items S such that he has a complete preference list over items from $I - S$ and he prefers to stay unmatched over an assignment of an item from S , and (b) relaxation of the unit-demand assumption: where an agent has demand only for subsets of size K and he has a complete preference list over all subsets of size K , for some parameter K . In both cases, we analyze the performance of RSD whose definition can be easily modified to both cases. We leave the extension and analysis of PS to these cases as an open direction.

4.1 Incomplete Preference Lists

We assume that there are n agents and m items where each agent a has an *ordered* preference list L_a on some arbitrary subset of $\{1, 2, \dots, m\}$. We note that the proof of Theorem 1 can be modified easily to show that the ordinal welfare factor of RSD is at least $1/2$ even in this case.

Now we focus on the linear utility model. In this model, if agent a is allocated the j^{th} item on the list, denoted by $L_a(j)$, it gets a utility of $\frac{|L_a|+1-j}{|L_a|}$. Our main result is that the linear welfare factor of RSD drops to $\tilde{\Theta}(1/n^{1/3})$. We now focus here on establishing almost tight upper and lower bounds on the linear welfare factor of RSD.

Linear Welfare Factor of RSD with partial lists is $\tilde{O}(1/n^{1/3})$:

We start by showing a family of instances that establish an upper bound of $\tilde{O}(1/n^{1/3})$ on the linear welfare factor of RSD. The set A of agents is partitioned into two sets, namely, a set $G = \{g_1, g_2, \dots, g_a\}$ of *good* agents, and a set $B = \{b_1, b_2, \dots, b_b\}$ of *bad* agents, where $a = n^{2/3}$ and

$b = n - n^{2/3}$. The set of items is partitioned into two sets $I_G = \{1, 2, \dots, p\}$ and $I_B = \{1', 2', \dots, q'\}$ where $p = n^{2/3}$ and $q = n^{1/3}$. For each $i \in [1..a]$, the preference list of agent g_i consists of a single item, namely, $\{i\}$. For each $j \in [1..b]$, the preference list of agent b_j consists of items in I_B ordered as $\{1', 2', \dots, q'\}$, followed by the item $i \in I_G$ such that $i = 1 + (j \bmod n^{2/3})$. An agent $b_j \in B$ whose preference list ends in an item $i \in I_G$, is referred to as a *type- i* bad agent; there are $(n^{1/3} - 1)$ type- i bad agents for each $i \in [1..a]$.

It is easy to verify that OPT is at least $n^{2/3}$, since any allocation that assigns items in I_G to agents in G in accordance with the unique matching defined by the preference lists of agents in G achieves $n^{2/3}$ utility. We now show that in contrast, the expected utility of RSD on this instance is bounded by $O(n^{1/3} \log n)$, yielding the desired lower bound.

Let $t_1 = 2n^{2/3}$ and let $t_2 = 9n^{2/3} \ln n$. The analysis relies on the probability of occurrence (over the choice of random arrival ordering of the agents) of the four events defined below. We bound the probabilities of these events by careful application of Chernoff's inequality.

\mathcal{E}_0 : We say that the event \mathcal{E}_0 occurs if there exists an $i \in [1..a]$ such that the number of type- i bad agents that arrive in the first t_1 steps exceeds $(n^{1/3} - 1)/2$. Then

$$\Pr[\mathcal{E}_0] \leq (n^{2/3})e^{-\Omega(n^{1/3})} = e^{-\Omega(n^{1/3})}$$

\mathcal{E}_1 : We say that the event \mathcal{E}_1 occurs if the number of good agents that arrive in the first t_1 steps exceeds $n^{1/3}$. Then $\Pr[\mathcal{E}_1] \leq e^{-\Omega(n^{1/3})}$.

\mathcal{E}_2 : We say that event \mathcal{E}_2 occurs if there exists an $i \in [1..a]$ such that no type- i bad agent arrives in the interval $(t_1, t_2]$. Then

$$\Pr[\mathcal{E}_2] = \Pr[\mathcal{E}_2 \mid \mathcal{E}_0] \cdot \Pr[\mathcal{E}_0] + \Pr[\mathcal{E}_2 \mid \overline{\mathcal{E}_0}] \cdot \Pr[\overline{\mathcal{E}_0}] \leq \Pr[\mathcal{E}_0] + \Pr[\mathcal{E}_2 \mid \overline{\mathcal{E}_0}] \leq 1/n^2$$

\mathcal{E}_3 : We say that the event \mathcal{E}_3 occurs if the total number of good agents that arrive in the interval $[1, t_2]$ exceeds $18n^{1/3} \ln n$. Then $\Pr[\mathcal{E}_3] \leq 1/n^2$.

Thus $\Pr[\mathcal{E}_1 \vee \mathcal{E}_2 \vee \mathcal{E}_3] \leq 3/n^2$.

We now analyze the performance of RSD when none of the events \mathcal{E}_1 , \mathcal{E}_2 , or \mathcal{E}_3 occur. If the event \mathcal{E}_1 does not occur, then by time t_1 , at least $n^{1/3}$ agents from B must have arrived, so by time t_1 , all items in I_B must have been assigned to agents in B . Total contribution from items in I_B is thus bounded by $n^{1/3}$. Additionally, if the event \mathcal{E}_2 does not occur, we are guaranteed that by time t_2 , all items in I_G must have been assigned to agents in $B \cup G$. Finally, if the event \mathcal{E}_3 also does not occur, then at most $18n^{1/3} \ln n$ items in I_G get assigned to agents in G , and the rest are assigned to agents in B . Thus the total contribution from items in I_G can be bounded by $O(n^{1/3} \log n)$. Hence the expected utility of RSD is bounded by

$$\Pr[\mathcal{E}_1 \vee \mathcal{E}_2 \vee \mathcal{E}_3] \cdot (n^{1/3} + n^{2/3}) + \Pr[\overline{(\mathcal{E}_1 \vee \mathcal{E}_2 \vee \mathcal{E}_3)}] \cdot O(n^{1/3} \log n) = O(n^{1/3} \log n).$$

This completes the proof that the linear welfare factor of RSD is at most $\tilde{O}(1/n^{1/3})$.

Linear Welfare Factor of RSD with incomplete preference lists is $\tilde{\Omega}(1/n^{1/3})$:

We now show that the linear welfare factor of RSD is at least $\tilde{\Omega}(1/n^{1/3})$. We start by making two easy observations, namely, (a) OPT is at most n , (b) linear welfare factor of RSD is $\max\{\Omega(1/\text{OPT}), \Omega(\text{OPT}/n)\}$. It thus suffices to show that linear welfare factor of RSD is $\tilde{\Omega}(1/n^{1/3})$ when $n^{1/3} \leq \text{OPT} \leq n^{2/3}$.

Let G and I_G respectively denote the subset of agents and items that contributes to OPT , that is, there exists a perfect matching M between G and I_G . Note that $|G| = |I_G| \geq \text{OPT}$. We start with the assumption that each agent in G contributes $\Omega(1)$ – this case essentially captures the main idea of the analysis. We then briefly sketch the slight modification needed to eliminate this assumption.

First observe that if RSD is run on the agents in G alone, then for a uniformly at random ordering of agents in G , in expectation, $\Omega(G)$ agents match to items along edges of utility $\Omega(1)$. Now when RSD is run on the entire set of agents, if the total expected utility achieved by RSD from agents in G is at least $\text{OPT}/(2n^{1/3})$, then we are already done. Otherwise, there must exist a set B_1 of *bad* agents (need not be disjoint from G) such that (a) the list of each agent in B_1 contains at least $2n^{1/3}$ elements, and (b) $|B_1| \geq 2n^{1/3}|G|$. Let

$$I_1 = \bigcup_{i \in B_1} \bigcup_{1 \leq j \leq n^{1/3}} L_i(j)$$

Observe that for *any* ordering of agents in B_1 , the first $n^{1/3}$ agents in B_1 are able to match to an item in the first-half of their respective lists, deriving a utility of at least $n^{1/3}/2$. Thus in absence of any other agents, any execution of RSD on agents in B_1 alone derives a total utility of $n^{1/3}/2$ or more. Now when RSD is run on the entire set of agents, if the total expected utility in RSD from agents in B_1 is at least $(n^{1/3})/2$, we are once again done since $\text{OPT} \leq n^{2/3}$ by our earlier assumption.

If not, there must exist a set B_2 of *very bad* agents (need not be disjoint from $G \cup B_1$) such that

$$\left(\frac{|B_1|}{|B_2|}\right) \left(\frac{n^{1/3}}{2}\right) < \frac{\text{OPT}}{2n^{1/3}}.$$

Thus

$$|B_2| > \frac{n^{2/3}|B_1|}{\text{OPT}} \geq 2n,$$

a contradiction!

Finally, to eliminate the assumption that OPT matches agents in G to the set I_G along edges of utility $\Omega(1)$, we note that there must exist a set $G' \subseteq G$ and a set $I' \subseteq I_G$ such that for some $\alpha \in (0, 1]$, (i) OPT matches items in I' to agents in I_G along edges of utility $\Theta(\alpha)$, and (ii) the utility contributed by this matching is $\Omega(\text{OPT}/\log n)$. This follows from a standard geometric grouping argument. We can now focus on agents-item pairs defined by G' and I' , and apply the same argument as above.

Without Randomness: We note that randomness is crucial to achieving non-trivial linear welfare factor in case of partial preference lists. If the permutation σ on agents is chosen deterministically, then the linear welfare factor drops to $\Theta(\log n/n)$. We refer to this algorithm as the *serial dictatorship* (SD) algorithm.

An $O(\log n/n)$ bound on the linear welfare factor of SD follows from the following family of instances where OPT is $n/2$ while there exists an arrival order on which SD achieves only $O(\log n)$ overall utility. Assume that n is an even integer. There are $n/2$ items in the instance, numbered 1 through $n/2$. The ordered list L_i is defined to be $\{1, 2, \dots, i\}$, for $1 \leq i \leq n/2$, and $L_i = \{i - (n/2)\}$, for $n/2 < i \leq n$. Suppose the agents arrive in order 1 through n . Then it is easy to see that for each $i \in \{1, 2, \dots, n/2\}$, agent i chooses item i and obtains utility of $1/i$. Thus the total utility

achieved is $H(n/2) = \Theta(\log n)$. On the other hand, the optimal allocation assigns item i to agent $i + (n/2)$, for $1 \leq i \leq n/2$, achieving an overall utility of $n/2$.

On the other hand, SD always achieves a linear welfare factor of $\Omega(\log n/n)$. Fix an optimal solution, say OPT, and let k be the number of items allocated in OPT. Clearly, $k \geq |\text{OPT}|$. Since SD always outputs a maximal matching between agents and items, it follows that the total number of items allocated in any execution of SD must be at least $k/2$. Fix an execution of SD, and let I' be the ordered set of items that get allocated in SD. Then it is easy to see that when the j^{th} item in I' is allocated, the utility derived is at least $1/j$. Thus the total utility obtained by SD is at least $H(k/2) = \Omega(\log k)$. The performance ratio of SD is at most $\max_{1 \leq k \leq n} k/H(k/2) = \Theta(n/\log n)$.

4.2 Non-Unit Demands

We now analyze the ordinal welfare factor of RSD when each agent desires a bundle of K items for some integer parameter K and each agent has a preference list that defines an ordering over the subsets of items of size K . We show that in this case, the ordinal welfare factor of RSD is $\Theta(1/K)$. We will show that an upper bound of $1/K$ holds on the RSD ordinal welfare factor even when preference lists are complete while the RSD ordinal welfare factor is at least $\Omega(1/K)$ if the preference lists are partial.

To see that the ordinal welfare factor is bounded from above by $1/K$, consider the following instance with n agents and nK items where the agents are partitioned into $p = n/K$ groups of K agents each. Let A_1, A_2, \dots, A_p denote these groups of agents. For each $1 \leq j \leq K$ and $1 \leq i \leq p$, the preference list of the j^{th} agent in group A_i is as follows: the first choice is the item set $\{(i-1)K^2 + 1, (i-1)K^2 + K + 1, \dots, (i-1)K^2 + (K-1)K + 1\}$, the second choice is the item set $\{(i-1)K^2 + (j-1)K + 1, (i-1)K^2 + (j-1)K + 2, \dots, (i-1)K^2 + (j-1)K + K\}$, the next $\binom{K^2}{K} - 2$ choices consist of an arbitrary ordering of the remaining K -element subsets of items $\{(i-1)K^2 + 1, (i-1)K^2 + 2, \dots, (i-1)K^2 + K^2\}$, and finally, this is followed by an arbitrary ordering on the remaining subsets. Thus the first choice of each agent in any group A_i intersects with the first and second choices of every other agent in the group. Moreover, it is easy to see that each agent also matches to one of the first $\binom{K^2}{K}$ choices on its list. Thus allocations made to agents in one group never interfere with the allocations made to agents in another group. Now fix the matching M that allocates to each agent its second preference – it is easy to verify that this allocation is conflict-free. Consider now a run of the RSD algorithm. The first agent to arrive in a group A_i always gets his first preference. However, once this allocation is done, no other agent in the group A_i can be allocated any item. Thus the ordinal welfare factor of RSD on this instance is precisely $1/K$.

On the other hand, an argument similar to that in Theorem 1 shows that the performance of RSD is at least $\frac{1}{2K} - o(\frac{1}{K})$. The argument is the following: by the time the t^{th} agent arrives, at most Kt items have been taken away. Of the $(n-t)$ agents remaining, at least $n - (K+1)t$ of them have their optimal bundle intact, as Kt items can hit at most the bundles of Kt agents. So, the probability the t^{th} agent is happy is at least $\frac{n-(K+1)t}{n-t}$. Therefore, the expected number of happy agents is at least

$$\sum_{t=1}^{n/(K+1)} (n - (K+1)t)/(n-t)$$

The expression simplifies to

$$n - Kn \cdot \sum_{t=1}^{n/(K+1)} \frac{1}{n-t} \approx n(1 - K(\ln(1 + 1/K))) \geq n/2K - o(1/K)$$

5 Conclusions and Open Directions

In this paper, we studied the social welfare of two well studied mechanisms, RSD and PS, for one-sided matching markets. We focussed on two measures: one was the ordinal welfare factor which we believe is a new way to measure efficiency when agents utilities are obtained as rankings or preference lists, the other was the natural linear utilities measure which assumed that the utility of an agent dropped linearly down his preference list.

We performed a tight analysis of the ordinal welfare factors of both mechanisms, and the linear welfare factor in the case of efficient instances: instances where the optimum matching gave every agent one of his first $o(n)$ top choices. An open problem is to perform a tight analysis of the case of linear welfare factor for general instances.

We think the notion of ordinal welfare factor and the definitions of truthfulness in case of ordinal utilities will be useful for other problems as well where the utilities are expressed as preference lists rather than precise numbers. Examples which come to mind are scheduling, voting, and ranking applications.

References

- [1] A. Abdulkadiroglu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66:689 – 701, 1999.
- [2] A. Abdulkadiroglu and T. Sönmez. Ordinal efficiency and dominated sets of assignments. *Journal of Economic Theory*, 112:157 – 172, 2003.
- [3] A. Abdulkadiroglu and T. Sönmez. Matching markets: Theory and practice. *prepared for the Econometric Society World Congress, China*, 2010, available at first author’s website.
- [4] D. Abraham, R. Irving, T. Kavitha, and K. Melhorn. Popular matchings. *SIAM Journ. of Computing*, 37:1030 – 1045, 2007.
- [5] G. Aggarwal, G. Goel, C. Karande, and A. Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings, ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.
- [6] I. Ashlagi, F. Fischer, I. Kash, and A. D. Procaccia. Mix and Match. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC’10)*, 2010.
- [7] B. Birnbaum and C. Mathieu. On-line bipartite matching made simple. *ACM SIGACT News*, 39, 2008.
- [8] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295 – 328, 2001.
- [9] A. Bogomolnaia and H. Moulin. A simple random assignment problem with a unique solution. *Economic Theory*, 19:623 – 635, 2002.

- [10] E. Budish, Y. K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *Unpublished Manuscript available at first author's personal webpage.*, October, 2010.
- [11] H. Cres and H. Moulin. Scheduling with opting out: Improving upon random priority. *Operations Research*, 49:565 – 577, 2001.
- [12] O. Dekel, F. A. Fischer, and A. D. Procaccia. Incentive compatible regression learning. *Journal of Computer and System Sciences*, 76(8), 2010.
- [13] S. Dughmi and A. Ghosh. Truthful assignment without money. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC'10)*, 2010.
- [14] D. Gale. College course assignments and optimal lotteries. *University of California, Berkeley Technical Report*, 1987.
- [15] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9 – 15, 1962.
- [16] P. Gardenfors. Match making: assignments based on bilateral preferences. *Behavioral Sciences*, 20:166 – 173, 1975.
- [17] G. Goel and A. Mehta. Online budgeted matching in random input models with applications to adwords. In *Proceedings, ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008.
- [18] M. Guo and V. Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-10)*, 2010.
- [19] M. Guo, V. Conitzer, and D. Reeves. Competitive repeated allocation without payments. In *Fifth Workshop on Internet and Network Economics (WINE-09)*, 2009.
- [20] Chien-Chung Huang, Telikepalli Kavitha, Dimitrios Michail, and Meghana Nasre. Bounded unpopularity matchings. In *Proceedings of the 11th Scandinavian workshop on Algorithm Theory*, 2008.
- [21] A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *American Mathematical Monthly*, 69:9 – 15, 1962.
- [22] R.M. Karp, U.V. Vazirani, and V.V. Vazirani. An optimal algorithm for online bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, 1990.
- [23] A. K. Katta and J. Sethurman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131:231 – 250, 2006.
- [24] T. Kavitha, J. Mestre, and M. Nasre. Popular mixed matchings. *Theoretical Computer Science*, Article in Press, 2010.
- [25] Richard Matthew McCutchen. The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In *Proceedings of the 8th Latin American conference on Theoretical informatics*, 2008.

- [26] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC'09)*, 2009.
- [27] A. E. Roth, T. Sönmez, and U. Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119:457 – 488, 2004.
- [28] A. E. Roth and M. Sotomayor. *Two-Sided Matching: A Study on Game- Theoretic Modelling*. Cambridge University Press, 1990.
- [29] L. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics.*, 1:23 – 28, 1974.
- [30] T. Sönmez and U. Ünver. Matching, allocation, and exchange of discrete resources. *Handbook of Social Economics*, forthcoming, available at second author's website.
- [31] L. Svensson. Strategyproof allocation of indivisible goods. *Social Choice and Welfare*, 16:557 – 567, 1999.
- [32] L. Zhou. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory*, 52:123 – 135, 1990.

A Comparison of PS and RSD

A question might arise whether one of the two algorithms is as good as the other in all instances, when one considers the two efficiency measures studied in the paper. In this subsection we show, via examples, that such a statement doesn't hold true for either measure. Both examples are taken from the paper of Bogomolnaia and Moulin [8].

Instance 1

			a	b	c	d			a	b	c	d
1	a	b	c	d			1	$5/12$	$1/12$	$5/12$	$1/12$	
2	a	b	c	d			2	$5/12$	$1/12$	$5/12$	$1/12$	
3	b	a	d	c			3	$1/12$	$5/12$	$1/12$	$5/12$	
4	b	a	d	c			4	$1/12$	$5/12$	$1/12$	$5/12$	

Instance 2

		a	b	c		a	b	c				
1	a	b	c		1	$1/2$	$1/6$	$1/3$	1	$1/2$	$1/4$	$1/4$
2	a	c	b		2	$1/2$	0	$1/2$	2	$1/2$	0	$1/2$
3	b	a	c		3	0	$5/6$	$1/6$	3	0	$3/4$	$1/4$

In the two instances above, the first table shows the preference lists of agents (numbered 1,2,...) over items (denoted as *a, b, ...*). The second table shows the randomized allocation obtained from the RSD mechanism, the third shows that obtained via the PS mechanism.

Let us consider the linear utilities case first. In the first instance, RSD gets a utility of $4 \cdot \frac{1}{4} \cdot (4 \cdot 5/12 + 3 \cdot 1/12 + 2 \cdot 5/12 + 1 \cdot 1/12) = 34/12$. While PS obtains a utility of $4 \cdot \frac{1}{4} \cdot (4 \cdot 1/2 + 2 \cdot 1/2) = 3$. So PS does better than RSD. In the second instance, the calculation shows

the utility of RSD is greater than that of PS: the utilities on the allocations where they differ for RSD is $\frac{1}{3} \cdot (2 \cdot 1/6 + 1 \cdot 1/3 + 3 \cdot 5/6 + 1 \cdot 1/6) = 10/9$, while for PS is $\frac{1}{3} \cdot (2 \cdot 1/4 + 1 \cdot 1/4 + 3 \cdot 3/4 + 1 \cdot 1/4) = 13/12$.

Moving to ordinal welfare factor, in the second instance, if the benchmark matching is $(1, a), (2, c), (3, b)$, then the expected number of agents made happy by PS is $9/4$, while that in RSD is $7/3$. However if the benchmark matching is $(1, b), (2, a), (3, c)$, then the ordinal welfare factors of PS and RSD are $9/4$ and $13/6$ respectively. It is easy to check that both benchmark matchings are pareto-optimal.

B Performance of RSD with linear utilities

Let α and β be two parameters between 0 and 1 to be fixed later. We assume that at least αn agents get an item which is in their top βn choices. It is easy to see that if not, the optimum is at most $(1 - \beta + \alpha\beta)n$ and RSD gives a linear welfare factor of $\frac{1}{2(1-\beta+\alpha\beta)}$.

Call these αn agents good. Now consider the argument in Lemma 3. At time t , the expected number of remaining good agents is $\alpha(n - t)$. Thus, with probability at least $\left(\alpha - \frac{d_t}{n-t}\right)$, we choose a good agent and he gets an item from his top βn choices, thus he gets a utility of at least $1 - \beta$, and so the benefit over the ‘trivial’ $(1 - t/n)$ utility is $(t/n - \beta)$. Thus, the total benefit obtained at time t is at least $\left(\alpha - \frac{d_t}{n-t}\right) \cdot (t/n - \beta)$. Using the fact that $d_t \leq (t+2)(n-t)/(n+1)$, we get that the total benefit over 50% is at least $\int_{t_1}^{t_2} \left(\alpha - \frac{t+2}{n+1}\right) \left(\frac{t}{n} - \beta\right)$, where $t_1 = \beta n$ (when the second term becomes positive), and $t_2 = \alpha * (n+1) - 2$ (after which the first term becomes negative). Therefore, the RSD mechanism gives an approximation of at least the minimum of this integral and $\frac{1}{2(1-\beta+\alpha\beta)}$, for any α and β . For $\alpha = 0.77$ and $\beta = 0.22$, we get the above expression to be 52.6%.

C Performance of PS with linear utilities

We now consider the performance of PS with linear utilities even when the instance may not be efficient, i.e. some agents get less than $1 - o(1)$ in the optimum matching. Let M^* be the optimal matching.

Consider items in the order in which they are fully allocated in the run of PS algorithm (break ties arbitrarily). For an item i in this order, let a be the agent such that $M^*(a) = i$ and let u_i be its utility for item i . The social optimum is $\sum_i u_i$. We denote it by **OPT**. Now we compute the total utility of items assigned to a in PS. From Observation ??, we get that in the time interval $[\frac{k-1}{n}, \frac{k}{n}]$ for $1 \leq k \leq i$, agent a will obtain utility at a rate of at least $\max\{u_i, (n - k + 1)/n\}$. This is because in this time interval both $M^*(a)$ and his k^{th} choice are available. Similarly, for $k > i$, agent a gets utility at a rate $(n - k + 1)/n$ in the interval $[\frac{k-1}{n}, \frac{k}{n}]$. So, the total utility of agent a is at least $\sum_{k=1}^n \frac{n-k+1}{n} \geq 1/2$ if $u_i \leq (n - i)/n$, and it is $\frac{1}{2} + \frac{(u_i - \frac{n-i}{n})^2}{2}$ otherwise. We denote the total utility of the assignment in the PS algorithm by **ALG**.

Observe that, if for any i , $u_i < (n - i)/n$, then increasing u_i increases the **OPT** but leaves **ALG** unchanged. Since we are looking at an upper bound on the performance PS algorithm, it is safe to assume that, for every i , $u_i \geq (n - i)/n$. We want to show that there exists a constant $\beta \leq 1$ such that

$$\text{ALG} \geq \beta \text{OPT} \tag{5}$$

for any set of values of $u_i, 0 \leq u_i \leq 1$. By taking a partial derivative w.r.t. u_i , and using the fact that $u_i \leq 1$, we see that the minimizer of **ALG** is $\min\left(1, \frac{n-i}{n} + \beta\right)$. Now we compute the values of

OPT and ALG with these values of u_i .

$$\text{OPT} = \sum_i (u_i) \tag{6}$$

$$= \sum_{1 \leq i \leq \beta n} 1 + \sum_{\beta n \leq i \leq n} \left(\beta - \frac{i}{n} \right) \tag{7}$$

$$= \frac{n}{2} + \beta n - \frac{\beta^2}{2} \tag{8}$$

$$\text{ALG} = \frac{n}{2} + \sum_{1 \leq i \leq n} \frac{\left(u_i - \frac{n-i}{n} \right)^2}{2} \tag{9}$$

$$= \frac{n}{2} + \sum_{1 \leq i \leq \beta n} \frac{i^2}{2} + \sum_{\beta n \leq i \leq n} \frac{\beta^2 n}{2} \tag{10}$$

$$= \frac{n}{2} + \frac{\beta n}{2} (1 - \beta) + \frac{\beta^3}{6} \tag{11}$$

Recall that β should satisfy $\text{ALG} \geq \beta \text{OPT}$. Solving for this gives $\beta = 0.6602$. It will be interesting to close the curiously small gap between the upper bound and the lower bound.